arXiv:1512.03246v1 [cs.CC] 10 Dec 2015

New Deterministic Algorithms for Solving Parity Games*

Matthias Mnich† Heiko Röglin† Clemens Rösner†

December 12, 2015

Abstract

We study parity games in which one of the two players controls only a small number k of nodes and the other player controls the $n-k$ other nodes of the game. Our main result is a fixed-parameter algorithm that solves bipartite parity games in time $k^{O(\sqrt{k})} \cdot O(n^3)$, and general parity games in time $(p+k)^{O(\sqrt{k})} \cdot O(pnm)$, where p is the number of distinct priorities and m is the number of edges. For all games with $k = o(n)$ this improves the previously fastest algorithm by Jurdziński, Paterson, and Zwick (SICOMP 2008).

We also obtain novel kernelization results and an improved deterministic algorithm for graphs with small average degree.

1 Introduction

A parity game [5] is a two-player game of perfect information played on a directed graph G by two players, *even* and *odd*, who move a token from node to node along the edges of G so that an infinite path is formed. The nodes of G are partitioned into two sets V_0 and V_1; the even player moves if the token is at a node in V_0 and the odd player moves if the token is at a node in V_1. The nodes of G are labeled by a *priority function* $p : V \to \mathbb{N}_0$, and the players compete for the parity of the highest priority occurring infinitely often on the infinite path $v_0, v_1, v_2 \dots$ describing a play: the even player wins if $\limsup_{i\to\infty} p(v_i)$ is even, and the odd player wins if it is odd.

The winner determination problem for parity games is the algorithmic problem to determine for a given parity game $G = (V_0 \uplus V_1, E, p)$ and an initial node $v_0 \in V_0 \cup V_1$, whether the even player has a winning strategy in the game if the token is initially placed on node v_0. We say that an algorithm for this problem *solves* parity games. Parity games have various applications in computer science and the theory of formal languages and automata in particular. They are closely related to other games of infinite duration, such as mean payoff games, discounted payoff games, and stochastic games [10]. Solving parity games is linear-time equivalent to the model checking problem for the modal μ-calculus [19]. Hence, any parity game solver is also a model checker for the μ-calculus (and vice versa).

Many algorithms have been suggested for solving parity games [4, 12, 20, 21], yet none of them is known to run in polynomial time. McNaughton [15] showed that the winner determination problem belongs to the class $\mathsf{NP} \cap \mathsf{coNP}$, and Jurdziński [10] strengthened this to $\mathsf{UP} \cap \mathsf{coUP}$. It is a long-standing open question whether parity games can be solved in polynomial time. The fastest known deterministic algorithm is due to Jurdziński, Paterson, and Zwick [12] and it has a run time of $n^{O(\sqrt{n})}$ for general parity games and of $n^{O(\sqrt{n/\log n})}$ for parity games in which every node has out-degree at most two. The fastest known randomized algorithm for general parity games is due to Björklund et al. [4] and it has a run time of $n^{O(\sqrt{n/\log n})}$.

As a polynomial-time algorithm for solving parity games has remained elusive, researchers have started to consider which restrictions on the game allow for polynomial-time algorithms. One such well-studied restriction is the treewidth t of the underlying undirected graph G of the game. Obdržálek [16] found an

*This research was supported by ERC Starting Grant 306465 (BeyondWorstCase).

†Department of Computer Science, University of Bonn, Germany. {mmnich@,roeglin@cs.,roesner@cs.}uni-bonn.de

algorithm solving parity games on n nodes in time $n^{O(t^2)}$. Later, Fearnley and Lachish [6] gave an algorithm solving parity games in time $n^{O(t \log n)}$. Another well-studied parameter for parity games is the number p of distinct priorities by which the nodes of the game are labeled. The progress-measure lifting algorithm by Jurdziński [11] solves parity games in time $O(pm(2n/p)^{p/2})$, where m denotes the number of edges of G. This run time has been improved by Schewe [18] to $O(m((2e)^{3/2}n/p)^{p/3})$. Fearnley and Schewe [7] presented an algorithm for solving parity games with run time $O(n(t+1)^{t+5}(p+1)^{3t+5})$, assuming that a tree decomposition of G with width t is given.

For a given parameter κ, one usually aims for *fixed-parameter algorithm* algorithms, i.e., algorithms that run in time $f(\kappa) \cdot n^c$ for some computable function f and some constant c that is independent of κ. Such an algorithm can be practical for large instances if f grows moderately and c is small. From the previously mentioned algorithms only the algorithm by Fearnley and Schewe [7] is a fixed-parameter algorithm for the combined parameter (t, p). It is not known if fixed-parameter algorithms exist for the parameter t or the parameter p alone.

Further parameters for which polynomial-time algorithms for parity games have been suggested include DAG-width [1], clique-width [17], and entanglement [3]; none of these are fixed-parameter algorithms.

1.1 Our Contributions

We study as parameter the number k of nodes that belong to the player who controls the smaller number of nodes in the parity game. Our first result is a *subexponential* fixed-parameter algorithm for solving general parity games for parameters p and k and for parameter only k for bipartite parity games (where players alternate between their moves).

Theorem 1. *There is a deterministic algorithm that solves any parity game G on n nodes and m edges in time $(p+k)^{O(\sqrt{k})} \cdot O(pnm)$, where k denotes the minimum number of nodes owned by one of the players and p the number of distinct priorities. If G is bipartite, the algorithm runs in time $k^{O(\sqrt{k})} \cdot O(n^3)$.*

Thus, our algorithm is particularly efficient if the game is unbalanced, in the sense that one player owns only k nodes and the other player owns the remaining $n - k \gg k$ nodes.

Let us remark that it is not very hard to show fixed-parameter tractability for parameter $p + k$; indeed McNaughton's algorithm [15] can be shown to run in time $p^k \cdot n^{O(1)}$, and this was improved to $p^{\log k} \cdot 4^k \cdot n^{O(1)}$ by Gajarský et al. [8]. Our key contribution here is to reduce the dependence of k to a *subexponential* function. Indeed, this improvement allows us to derive the following immediate corollary of Theorem 1 to expedite the run time for solving *general* parity games.

Corollary 1. *There is a deterministic algorithm that solves parity games in time $n^{O(\sqrt{k})}$.*

Our algorithm is asymptotically always at least as fast as the fastest known deterministic parity game solver by Jurdziński, Paterson, and Zwick [12], which runs in time $n^{O(\sqrt{n})}$. For the case $k = o(n)$, our algorithm is asymptotically faster than theirs and constitutes the fastest known deterministic solver for such games.

We also prove the existence of a small kernel, as our second result. For a parameterized problem, a *kernelization algorithm* takes as input an instance x with parameter κ and computes in time $(|x| + \kappa)^{O(1)}$ an equivalent instance x' with parameter κ' (a *kernel*) with size $|x'| \leq g(\kappa)$, for some computable function g; here, equivalent means that an optimal solution for x can be derived in polynomial time from an optimal solution of x'.

Theorem 2. *Parity games can be kernelized in time $O(pmn)$ to at most $(p+1)^k + (p+1)k$ nodes, and bipartite parity games can be kernelized in time $O(n^3)$ to at most $k + 2^k \cdot \min\{k, p\}$ nodes and at most $k2^k \cdot \min\{k, p\}$ edges.*

This kernelization result is not only interesting for its own sake, but it is also an important ingredient in the proof of Theorem 1.

As our third result, we generalize the algorithm by Jurdziński, Paterson, and Zwick [12] for parity games with maximum out-degree 2 to arbitrary out-degree Δ.

Theorem 3. *There is a deterministic algorithm that solves parity games on n nodes out of which s_j nodes have out-degree at most j in time*

$$n^{O\left(\min_{1\le j\le n}\left\{\sqrt{n-s_j}+\sqrt{\frac{s_j}{\log_j s_j}}\right\}\right)}.$$

Corollary 2. *There is a deterministic algorithm that solves parity games on n nodes with maximum out-degree Δ in time $n^{O(\sqrt{\log(\Delta)\cdot n/\log(n)})}$ and parity games on n nodes with average out-degree Δ in time $n^{O(\sqrt{\log(\log(n)\Delta)\cdot n/\log(n)})}$.*

1.2 Detailed Comparison with Previous Work

Let us discuss in detail how our results compare to previous work. It is well-known (cf. [14, Lemma 3.2]) and easy to prove that the treewidth of a complete bipartite graph equals the size of the smaller side. Since the treewidth of a graph can only decrease when deleting edges, the graph underlying a bipartite parity game in which one player owns k nodes has a treewidth of at most k. However, as it is not known if there exists a fixed-parameter algorithm for parameter treewidth, the result in Theorem 1 for the bipartite case does not follow from previous work about parity games with bounded treewidth. As a parity game in which one player owns k nodes can have up to n different priorities, also the fixed-parameter algorithm for the combined parameter (t, p) by Fearnley and Schewe [7] does not imply our result.

The algorithm of Jurdziński, Paterson, and Zwick [12] for parity games with maximum out-degree two with run time $n^{O(\sqrt{n/\log n})}$ can easily be generalized to arbitrary parity games at the expense of its run time. For this, one only needs to observe that every parity game can be transformed into a game with maximum out-degree two by replacing each node with a higher out-degree by an appropriate binary tree. This transformation increases the number of nodes from n to $\Theta(m)$ where m denotes the number of edges in the original parity game. Hence, the run time becomes $m^{O(\sqrt{m/\log m})} = n^{O(\sqrt{m/\log n})}$. For graphs with average out-degree $\Delta = \omega(\log\log n)$ the resulting run time of $n^{O(\sqrt{\Delta n/\log n})}$ is asymptotically worse than the run time we obtain in Corollary 2

For graphs in which the variance of the out-degrees is large, our algorithm can even be better than stated in Corollary 2. If, for example, there are $n^{1-\varepsilon}$ nodes with an arbitrary out-degree for some $\varepsilon > 0$ and all remaining nodes have constant out-degree at most c then our algorithm has a run time of $n^{O(\sqrt{\frac{n}{\log n}})}$ (the minimum in Theorem 3 is assumed for $j = c$). This matches the best known bound for randomized algorithms.

Gajarský et al. [8] present an algorithm that solves parity games in time $w^{O(\sqrt{w})} \cdot n^{O(1)}$, where w denotes the modular width of G. Since the modular width of a bipartite graph can be exponential in the size of the smaller side, Theorem 1 does not follow from this result.

2 Fundamental Properties of Parity Games

A parity game $G = (V_0 \uplus V_1, E, p)$ consists of a directed graph $(V_0 \uplus V_1, E)$, where V_0 is the set of *even* nodes and V_1 is the set of *odd* nodes, and a priority function $p : V_0 \cup V_1 \to \mathbb{N}_0$. We often abuse notation and also refer to $(V_0 \uplus V_1, E)$ as the graph G. For each node $v \in V(G)$, we denote by $N_G^+(v)$ and $N_G^-(v)$ the set of out-neighbors and in-neighbors of v in G, respectively.

Two standard assumptions about parity games are (1) that G is bipartite with $E \subseteq (V_0 \times V_1) \cup (V_1 \times V_0)$, and (2) that each node $u \in V$ has at least one outgoing edge $(u, v) \in E$. The first assumption is often made because it is easy to transform a non-bipartite instance into a bipartite instance. However, the usual transformation increases the number of nodes in V_i by an amount of $|\{v \in V_{1-i} \mid N_G^-(v) \cap V_{1-i} \neq \emptyset\}|$, and can therefore increase the parameter $k = \min\{|V_0|, |V_1|\}$ significantly. We therefore consider bipartite and non-bipartite instances separately in Theorem 1.

We write $n = |V(G)|$, $m = |E|$ and $p = |\{p(v) \mid v \in V(G)\}|$. The game is played by two players, the *even* player (or player 0) and the *odd* player (or player 1). The game starts at some node $v_0 \in V(G)$. The players construct an infinite path (a *play*) as follows. Let u be the last node added so far to the path. If $u \in V_0$, then player 0 chooses an edge $(u, v) \in E$. Otherwise, if $u \in V_1$, then player 1 chooses an edge $(u, v) \in E$. In either case, node v is added to the path and a new edge is then chosen by either player 0 or player 1. As each node has at least one outgoing edge, the path constructed can always be continued. Let $v_0, v_1, v_2, \ldots$ be the infinite path constructed by the two players and let $p(v_0), p(v_1), p(v_2), \ldots$ be the sequence of the priorities of the nodes on the path. Player 0 *wins* the game if the largest priority seen infinitely often is even, and player 1 wins if the largest priority seen infinitely often is odd.

We will define $p_1(v)$ as $p(v)$ if $p(v)$ is odd and as $-p(v)$ if $p(v)$ is even. This allows us to say that, in case $p_1(v) > p_1(u)$ for some $v, u \in V$, player 1 *prefers* $p(v)$ over $p(u)$. Observe that removing an arbitrary finite prefix of a play in a parity game does not change the winner; we refer to this property of parity games as *prefix independence*. A *strategy* for player $i \in \{0, 1\}$ in a game G specifies, for every finite path $v_0, v_1, \ldots, v_k$ in G that ends in a node $v_k \in V_i$, an edge $(v_k, v_{k+1}) \in E$. A strategy is *positional* if the edge $(v_k, v_{k+1}) \in E$ chosen depends only on the last node v_k visited and is independent of the prefix path $v_0, v_1, \ldots, v_{k-1}$. A strategy for player $i \in \{0, 1\}$ is *winning* (for player i) from a start node v_0 if following this strategy ensures that player i wins the game, regardless of which strategy is used by the other player.

The fundamental determinacy theorem for parity games [5, 9] says that for every parity game G and every start node v_0, either player 0 has a winning strategy or player 1 has a winning strategy. Furthermore, if a player has a winning strategy from some node in a parity game, then she also has a winning positional strategy from this node. From now on we will therefore, unless stated differently, assume every strategy to be positional. Given positional strategies s_0 on V_0 and s_1 on V_1 and a start node $v_0 \in V$ the infinite path starting in v_0 corresponding to these strategies consists of a finite prefix and an infinite recurrence of a cycle $C = C(s_0, s_1, v_0)$. We call C the cycle *corresponding to* s_0, s_1, v_0 and say that s_0 and s_1 *create* C. The parity of the highest priority $p(u)$ of all nodes $u \in V(C)$ in cycle C then determines the winner of the game. The *winning set* of player $i \in \{0, 1\}$ is the set $\mathsf{win}_i(G) \subseteq V$ of nodes of the game G from which player i has a winning strategy.

For $i \in \{0, 1\}$, an *i-dominion* is a set of nodes $D \subseteq V$ so that player i can win from every node of D, without leaving D and without allowing the other player to leave D. An example of an i-dominion is the set $\mathsf{win}_i(G)$, but there may be smaller subsets of $\mathsf{win}_i(G)$ that are i-dominions as well. Although finding i-dominions can be just as hard as finding $\mathsf{win}_i(G)$, searching only for dominions with certain properties (e.g. small dominions) can be easier. In our algorithm we will use the fact that once an i-dominion is found, it can easily be removed from the graph, leaving a smaller game to be solved.

Next, we recall some well-known results about parity games that form the basis of the algorithms for solving parity games by McNaughton [15] and Zielonka [21]. We include them here as our algorithm relies on them as well; for a detailed exposition we refer to Grädel et al. [9]. Fix a parity game $G = (V_0 \uplus V_1, E, p)$.

For $i \in \{0, 1\}$, a set $B \subseteq V(G)$ is *i-closed* if for every $u \in B$ the following holds (we use the notation $\neg i$ for the element $1 - i \in \{0, 1\}$):

- If $u \in V_i$, then there exists some $(u, v) \in E$ such that $v \in B$; and
- if $u \in V_{\neg i}$, then for every $(u, v) \in E$, we have $v \in B$.

In other words, a set B is i-closed if player i can always choose to stay in B while simultaneously player $\neg i$ cannot escape from it, i.e., B is a "trap" for player $\neg i$.

Lemma 1. *For each $i \in \{0, 1\}$, the set $\mathsf{win}_i(G)$ is i-closed.*

Let $A \subseteq V(G)$ be a set of nodes and let $i \in \{0, 1\}$. The *i-reachability set* of A is the set $\mathsf{reach}_i(A)$ of nodes in A together with all nodes in $V(G) \setminus A$ from which player i has a strategy σ to enter A at least once (regardless of the strategy of the other player); we call such a strategy σ an *i-reachability strategy* to A.

Lemma 2. *For $A \subseteq V(G)$ and $i \in \{0, 1\}$, the set $V(G) \setminus \mathsf{reach}_i(A)$ is $(\neg i)$-closed.*

We will from now on assume that the graph of the parity game we operate on is encoded as an adjacency list.

Lemma 3. *For every set $A \subseteq V(G)$ and $i \in \{0,1\}$, the set $\mathsf{reach}_i(A)$ can be computed in $O(m)$ time, where $m = |E|$ is the number of edges in the game.*

If $B \subseteq V(G)$ is such that for each node $u \in V(G) \setminus B$ there is an edge (u,v) with $v \in V(G) \setminus B$, then the sub-game $G - B$ is the game obtained from G by removing the nodes of B. We will only be using B's for which $V(G) \setminus B$ is an i-closed set for some i. In this case every node in $v \in V(G) \setminus B$ has at least one out-going edge (v,w) with $w \in V(G) \setminus B$ and $G - B$ will therefore be well-defined. The next lemmas show some useful properties of sub-games.

Lemma 4. *Let G' be a sub-game of G and let $i \in \{0,1\}$. If the node set of G' is i-closed in G, then $\mathsf{win}_i(G') \subseteq \mathsf{win}_i(G)$.*

The next lemma shows that if we know some non-empty subset U of the winning set of some player $\neg i$ in a game G, then computing the winning sets of both players in G can be reduced to computing their winning sets in the smaller game $G - \mathsf{reach}_{\neg i}(U)$.

Lemma 5. *For any parity game G and $i \in \{0,1\}$, if $U \subseteq \mathsf{win}_{\neg i}(G)$ and $U^* = \mathsf{reach}_{\neg i}(U)$, then $\mathsf{win}_{\neg i}(G) = U^* \cup \mathsf{win}_{\neg i}(G - U^*)$ and $\mathsf{win}_i(G) = \mathsf{win}_i(G - U^*)$.*

The next lemma complements Lemma 5 by providing a way to find a non-empty subset of the winning set of player $\neg i$ in a parity game G or to conclude that player i can win from every node in G.

Lemma 6. *Let G be a parity game with largest priority $p_{\max}$ and let $V_{p_{\max}} \subseteq V(G)$ be the set of nodes with priority $p_{\max}$. Let $i = p_{\max} \pmod 2$ and let $G' = G - \mathsf{reach}_i(V_{p_{\max}})$. Then $\mathsf{win}_{\neg i}(G') \subseteq \mathsf{win}_{\neg i}(G)$. Also, if $\mathsf{win}_{\neg i}(G') = \emptyset$, then $\mathsf{win}_i(G) = V$, i.e., player i wins from every node of G.*

3 Kernelization of Parity Games

In this section, we describe some reduction rules for parity games. Theses rules are such that we can efficiently compute the winning sets of the original parity game once we know the winning sets of the reduced game.

3.1 General Parity Games

Lemma 7. *Any parity game $G = (V_0 \uplus V_1, E, p)$ can be transformed in time $O(pmn)$ to a parity game $G' = (V_0' \uplus V_1', E', p')$ with $V_1' \subseteq V_1$ such that*

- *there are no edges inside V_1', and*
- *for each node $v \in V_0'$ either $N_G^+(v) \subseteq V_1'$ or $N_G^-(v) \subseteq V_1'$, and*
- *$|V_0'| \leq \min\{n + pk, (p+1)^k + pk\}$, where $k = |V_1|$.*

Moreover, G and G' have the same winning sets on V_1' and the winner of the remaining nodes of G can be computed either during the transformation or from the winning sets of G' in linear time.

Proof. We will modify G in multiple steps. We will slightly abuse notation and refer in every step to the parity game that we obtained in the step before as $G = (V_0 \uplus V_1, E, p)$. First we eliminate all edges inside V_1. This can easily be achieved by adding for each edge $e = (v,w) \in E$ with $v, w \in V_1$ a new node v_e with $p'(v_e) = p(w)$ to V_0 and by replacing the edge e by the two edges (v, v_e) and (v_e, w). Since the new node v_e has only a single outgoing edge, this transformation does neither change the winning sets nor the winning strategies.

Next, we remove certain cycles inside V_0 from the game. Let $W_0 \subseteq V_0$ denote all nodes in V_0 that are part of at least one cycle that lies completely inside V_0 and whose highest priority is even. Clearly player 0

can win from all nodes in $\mathsf{reach}_0(W_0)$ by enforcing that such a cycle is entered and never left again. Hence, we can remove $\mathsf{reach}_0(W_0)$ from the game according to Lemma 5. Let $W_1 \subseteq V_0$ denote all nodes that are left in V_0 and from which player 0 cannot reach V_1. Then all paths that start in some node $u \in W_1$ must end in some cycle that is completely contained in V_0. Since we have removed all cycles whose highest priority is even, the maximum priority of this cycle must be odd. Thus, player 1 wins from all nodes in $\mathsf{reach}_1(W_1)$. Hence, we can also remove $\mathsf{reach}_1(W_1)$ from the game according to Lemma 5.

We use again the notation $G = (V_0 \uplus V_1, E, p)$ to refer to the parity game obtained after the previously discussed steps. Since we have removed all cycles from V_0 whose highest priority is even, player 0 loses for sure if she does not leave V_0. Hence, we can assume without loss of generality that the play leaves V_0 from every starting node if player 0 plays an optimal strategy. Then for every node $v \in V_0$ player 0 uses a (possibly empty) path inside V_0 followed by an edge that leads to some node $w \in V_1$. To determine the winning sets of a strategy of player 0 it is not important to know the exact paths player 0 chooses. Rather, it suffices to know for each $v \in V_0$ which node $w \in V_1$ will be reached and what the highest priority on the chosen v-w-path is. To get rid of long paths, we add $p \cdot |V_1|$ new nodes to V_0, one node $v(p', w)$ for each pair of a priority p' and a node $w \in V_1$. Node $v(p', w)$ has a priority p' and its only out-neighbor is w. The winner does not change if player 0 goes from $v \in V_0$ directly to $v(p', w)$ and from there directly to $w \in V_1$ instead of taking some other path from v inside V_0 with maximum priority p', followed by an edge that leads to w. For all such paths we add the corresponding edge $(v, v(p', w))$ and can therefore delete all edges inside V_0 that do not end in one of the new nodes $v(p', w)$ without changing the winning sets of the game. Observe that this ensures that all out-neighbors of the new nodes $v(p', w)$ belong to V_1 while all in-neighbors of the old nodes $v \in V_0$ belong to V_1.

It can be the case that for some pair $(v, w) \in V_0 \times V_1$ there are multiple nodes $v(p', w)$ that can be reached from v. We can assume without loss of generality that if player 0 decides to go from v to w via one of these nodes then she chooses the one that is best for her, i.e., the one with lowest p_1-value. All edges from v to other nodes $v(p', w)$ can be removed.

The inequality $|V_0'| \leq n + \min\{m, k^2\} + pk$ follows directly from the previously discussed construction: initially V_0 consists of $n - k \leq n$ nodes, there are at most $\min\{m, k^2\}$ edges inside V_1 for which we create a new node v_e, and there are only pk new nodes $v(p', w)$. To get rid of the term $\min\{m, k^2\}$ we can identify each node v_e, which derived from an edge $e = (v, w)$ inside V_1, with the node $v(p(w), w)$. This ensures that there are only pk new nodes. To show that $|V_0'| \leq (p+1)^k + pk$ we can reduce the number of old nodes in V_1 to ensure that at most $(p+1)^k$ remain. At first we remove all nodes $v \in V_0$ with $N_G^-(v) = \emptyset$, because they obviously cannot be part of a cycle and we can compute in linear time to which winning set they belong, once we know to which winning set their out-neighbors belong. Now let v and v' be two such nodes in V_0 with $N_G^+(v) = N_G^+(v')$. We then identify v and v' without changing the winning sets in V_1, since all nodes in $N_G^+(v)$ must have a priority at least as high as $\max\{p(v), p(v')\}$. This is because the priority of any node in $N_G^+(v)$ ($N_G^+(v')$) corresponds to the highest priority on a path that starts in v (v') and therefore must be at least $p(v)$ ($p(v')$). Afterwards there remains at most one node $v \in V_0$ for each possible set $N_G^+(v)$.

Since $N_G^+(v)$ can contain at most one new node corresponding to w for each $w \in V_1$ and there are p different ones to choose from there are at most $(p+1)^k$ different possibilities for $N_G^+(v)$.

It remains to analyze the run time of the transformation. We consider the different steps of the reduction separately. The first step of removing all edges inside V_1 can be performed in $O(m)$ because we only need to check for every edge $e = (v, w) \in E$ if $v, w \in V_1$ and then remove one edge and add two edges and a node. The second step of removing dominions completely inside V_0 can be executed in time $O(\log(p) \cdot m)$ as follows. First, we solve the solitary game on $V_0 \setminus \mathsf{reach}_1(V_1)$ and remove the 0-reachability-set of its 0-winning set; this can be done in time $O(\log p \cdot m)$ [2]. Thereafter, we compute the 1-reachability set of $V_0 \setminus \mathsf{reach}_0(V_1)$ and remove it; this can be done in time $O(m)$ [12]. The third step of removing long paths inside V_0 can be performed as follows. The algorithm computes the best priority for player 0 that a path in V_0 from a node $v \in V_0$ to a node $w \in V_1$ can have. To determine which nodes in $v \in V_0$ can reach which nodes $w \in V_1$ via a path in V_0 whose highest priority has fixed value of p', consider the subgraph $G^{\leq p'}$ of G that is induced by the set $V^{\leq p'}$ of nodes of priority at most p' and remove from it those edges that start in V_1. We then consider the set of nodes with priority exactly p' and compute by DFS in time $O(m)$ all nodes in V_1 reachable from

them. Then we compute with DFS for each node in V_0 which of the nodes with priority p' they can reach. This takes a total of at most $2n$ applications of DFS for each priority and therefore in total $O(pmn)$ time. In the last step where we remove and contract some of the nodes in V_0 we can find all nodes without incoming edges in time $O(m)$ and we can order all remaining nodes by their outgoing edges in time $O(|V_1| \cdot (n+p+1))$ using a version of radix-sort, where we view the set of out-neighbors as an $(p+1)$-adic number with $|V_1|$ digits. Thereafter, in linear time we identify sets of nodes with the same outgoing neighbors and identify them in total time $O(n+m)$. □

3.2 Bipartite Parity Games

In this section we give some reduction rules that efficiently reduce any bipartite game $G = (V_0 \uplus V_1, E, p)$ to a structurally simpler bipartite game $G' = (V_0' \uplus V_1', E', p')$, such that the winning sets of G can be efficiently recovered from the winning sets of G'. After exhaustive application of the reduction rules, the reduced game G' will have size bounded by some function of k and p only, independent of the size of G.

The digraphs of our underlying parity game may have self-loops and bidirected edges, but (without loss of generality) no parallel edges between the same two nodes. Thus, whenever parallel edges arise during the application of one of the reduction rules, we remove one of them without explicit mention.

Lemma 8. *Let $G = (V_0 \uplus V_1, E, p)$ be a bipartite parity game, and let $u, v \in V_0$ be such that $N_G^+(v) \subseteq N_G^+(u)$ and $p_1(v) \geq p_1(u)$. Let G' be the parity game obtained from G by deleting the edges $\{(w,u) \in E \mid (w,v) \in E\}$. Then the winning sets of G and G' are equal.*

Proof. We show that an edge $(w',u) \in \{(w,u) \in E \mid (w,v) \in E\}$ can only be part of a winning strategy for player 1 on node w' if the edge (w',v) is part of a winning strategy for player 1 on w' as well. Therefore, after deleting (w',u), player 1 wins from w' in G' if and only if he wins from w' in G. Deleting the edges in $\{(w,u) \in E \mid (w,v) \in E\}$ does therefore not change the winning sets.

Assume that player 1 has a winning strategy $s_1 : V_1 \to V_0$ for w' with $s_1(w') = u$. Let $s_1' : V_1 \to V_0$ be defined by $s_1'(w') = v$ and $s_1'(w) = s_1(w)$ for all $w \in V_1 \setminus \{w'\}$. We claim that s_1' is a winning strategy for player 1 on w' as well. Assume that there exists a counter strategy s_0' for s_1' such that player 0 wins the game with starting node w'. We will define a strategy s_0 for player 0 and show that s_0 is a counter strategy for s_1. Note that s_0 will not necessarily be a positional strategy. For all $w \in V_0 \setminus \{u\}$, s_0 chooses the same successor as s_0', but on u it might change its behavior. Each time the play encounters u directly after encountering w', strategy s_0 chooses $s_0'(v)$ as the successor of u. Every other time the play encounters u, strategy s_0 chooses $s_0'(u)$ as the successor of u.

The play defined by s_0' and s_1' can then be transformed into the play defined by s_0 and s_1 by replacing every appearance of the sequence $w', v, s_0'(v)$ with the sequence $w', u, s_0'(v)$. Let C' be the cycle created by s_0' and s_1'; then C' is a winning cycle for player 0. Then s_0 and s_1 will also create the cycle C', if C' does not contain the sequence $w', v, s_0'(v)$. Let us therefore assume that C' contains the sequence $w', v, s_0'(v)$. Let C be the closed walk obtained, when replacing v with u in C'. After a finite prefix the play defined by s_0 and s_1 will consist of an infinite recurrence of C. Since we have $p_1(v) \geq p_1(u)$ player 0 is wining the play defined by s_0 and s_1. This contradicts that s_1 is a winning strategy for player 1. □

Lemma 9. *Let $G = (V_0 \uplus V_1, E, p)$ be a bipartite parity game, and let $u, v \in V_0$ be nodes with $N_G^+(u) = N_G^+(v)$ and $p(v) = p(u)$. Let G' be the parity game obtained from G by contracting u and v into a new node v' with priority $p(v)$. Then u and v belong to the same winning set $\mathsf{win}_i(G)$ in G and v' belongs to the winning set $\mathsf{win}_i(G')$ of the same player in G'. For all other nodes the winning sets of G and G' coincide.*

Proof. Note that u and v belong to the winning set of the same player i in G. We can assume that player 0 chooses the same successor for u and v in her optimal strategy. Then no cycle created by optimal strategies contains both u and v and, after the contraction, each simple cycle that does not contain both u and v is again a simple cycle with the same priorities. Also, each cycle in the contracted game either exists in the original game (i.e., it does not contain v') or an equivalent cycle, which can be created by replacing v' with v or u, exists in the original game. We can also map winning strategies in the original game, where u and v

have the same successor into winning strategies in the resulting game and vice versa by simply identifying the successors of v' with the successor of u and v and vice versa, while keeping the rest of the strategy. We again assume that in the winning strategies in the original game, v and u have the same successor w. We then set the successor of v' to w and set the successor of any node w' with successor v or u to v'. The other way around v and u get the same successor as v' and any node w' with successor v' gets either v or u as its successor, depending on which of the edges (w', v) and (w', u) exists in the original game. A pair of strategies and the pair of strategies, to which they are mapped to, then create corresponding cycles and must therefore either both be winning for player 1 or both be winning for player 0. □

Lemma 10. *Let $G = (V_0 \uplus V_1, E, p)$ be a bipartite parity game, and let $v \in V(G)$ be such that $N_G^-(v) = \emptyset$. Then for the parity game $G' = G - v$ and for $i \in \{0, 1\}$, any node $v' \neq v$ is winning for player i in G if and only if it is winning for player i in G'.*

Proof. The condition $N_G^-(v) = \emptyset$ implies that v cannot be part of any cycle. Let $v \in V_j$; then $v \in \mathsf{win}_j(G)$ is equivalent to the existence of some node $w \in \mathsf{win}_j(G) \cap N_G^+(v)$. Since all possible strategies for all nodes except v are also possible strategies in $G - \{v\}$, all nodes in $V \setminus \{v\}$ belong to the same winning set in G and in $G - v$. (Notice that $G - \{v\}$ is again a parity game.) Once we computed the winning sets for $G - \{v\}$, we can check in time $O(n)$ whether $v \in \mathsf{win}_j(G)$ or $v \in \mathsf{win}_{\neg j}(G)$. □

Lemma 11. *Let $G = (V_0 \uplus V_1, E, p)$ be a parity game with largest priority $p_{\max} = \max\{p(v) \mid v \in V(G)\}$. If $p^{-1}(z) = \emptyset$ for some $z \in \{1, \ldots, p_{\max}\}$ then let $G' = (V_0 \uplus V_1, E, p')$ be the parity game obtained from G by setting $p'(v) = p(v) - 2$ for all $v \in V$ with $p(v) > z$ and $p'(v) = p(v)$ for all $v \in V$ with $p(v) < z$. Then the winning sets of the games G and G' coincide.*

Proof. Let s_0 and s_1 be strategies for player 0 and player 1, respectively, and let $C = (v_0, v_1, \ldots, v_\ell)$ be the cycle created by these strategies when the game starts at some node v. The parity j of the largest element in the set $Q = \{p(v_0), \ldots, p(v_\ell)\}$ determines which player wins in G and the parity j' of the largest element in the set $Q' = \{p'(v_0), \ldots, p'(v_\ell)\}$ determines which player wins in G'. It is easy to see that our reduction ensures that $j = j'$. Since this is true for any cycle, the lemma follows. □

Corollary 3. *In any parity game with maximum priority $p_{\max}$ to which the reduction rule described in Lemma 11 cannot be applied anymore, the set of priorities is either $\{0, 1, \ldots, p_{\max}\}$ or $\{1, \ldots, p_{\max}\}$.*

Lemma 12. *Let $G = (V_0 \uplus V_1, E, p)$ be a bipartite parity game with $|V_1| = k$ that is reduced according to Lemmas 8–10. Then $|V_0| \leq 2^k \cdot \min\{k, p\}$.*

Proof. For each node $v \in V_0$ there are 2^k possible choices for $N_G^+(v)$. Lemma 8 yields that for two nodes $v \neq u \in V_0$ with $N_G^+(v) = N_G^+(u)$ we must have $N_G^-(v) \cap N_G^-(u) = \emptyset$. Lemma 10 then yields that there can be at most k nodes in V_0 for every possible choice of $N_G^+(v)$. Also Lemma 9 yields that for each possible choice of $(N_G^+(v), p(v))$ there exists at most one node in V_0. □

Lemma 13. *There exists a sequence of applications of the reduction rules described in Lemmas 8–11 with a total run time of $O(n^3)$ that leads to a game in which none of these rules applies anymore.*

Proof. We show for each reduction rule separately how to apply it exhaustively in time $O(n^3)$. Although it can happen, that some reductions corresponding to one of the rules could lead to allowing some other reductions which were not allowed before. Therefore we cannot only apply all reductions corresponding to one rule after all reductions corresponding to another rule.

Most of the run time will be necessary to test if Lemma 8 or Lemma 9 applies to an ordered pair of nodes. We will argue how to apply the reductions such that we do not have to test the same ordered pair of nodes more than once, yielding a total run time of $O(n^3)$.

To apply all reductions of Lemma 11, we first sort the nodes in increasing order of their priorities and create an order of subsets each containing all nodes with the same priority; this can be done in $O(n \log(n))$ time. We then save for each of the subsets if its corresponding priority is odd or even and unite consecutive sets with the same parity. If the parity of the priority in the first subset is even, all nodes in the i-th subset

get priority $i-1$; otherwise all nodes in the i-th subset get priority i. Uniting sets can be done in linear time, and we cannot unite more than n times. The time for applying Lemma 11 is thus $O(n^2)$.

To apply all reductions for Lemma 8, we need to check for each pair of nodes $\{u,v\} \subseteq V_0$ with $p_1(v) \geq p_1(u)$ whether $N_G^+(v) \subseteq N_G^+(u)$, and find all nodes w with $(w,u) \in E$ and $(w,v) \in E$. There are $O(n^2)$ node pairs $\{u,v\} \in V_0$ with $p_1(v) \geq p_1(u)$, which can easily be found using the order of subsets created for Lemma 11. Checking if $N_G^+(v) \subseteq N_G^+(u)$ and finding all nodes w with $(w,u) \in E$ and $(w,v) \in E$ can be done in time $O(n)$. The total run time for Lemma 8 therefore is $O(n^3)$.

To apply all reductions for Lemma 9 we need to check for each pair of nodes $\{u,v\} \subseteq V_0$ with $p(v) = p(u)$ whether $N_G^+(v) = N_G^+(u)$. There are $O(n^2)$ such pairs $\{u,v\}$ with $p(v) = p(u)$, which can easily be found using the order of subsets created for Lemma 11. Testing whether $N_G^+(v) = N_G^+(u)$ and identifying u and v can be done in time $O(n)$. The total run time for Lemma 9 therefore is $O(n^3)$.

To apply all reductions for Lemma 10, we only need to check for each node if it has incoming edges and possibly delete it. Testing a node can be done in constant time, and deleting a node takes at most linear time. The time for applying Lemma 10 is thus $O(n^2)$.

We will first apply all feasible reductions for Lemma 11, then all feasible reductions for Lemmas 8, 9 and 10. Any reduction that is now possible was not feasible in the beginning.

Observe that some reductions can result in other reductions becoming feasible. Since we do not change the out-neighborhood of any node in V_0, reductions corresponding to Lemmas 8 and 9 for a pair of nodes $\{u,v\} \subseteq V_0$ can only become feasible when we combine the two subsets containing v and u in a reduction corresponding to Lemma 11. For each node pair $\{u,v\} \subseteq V_0$ this can happen at most once. The total run time for all reductions corresponding to Lemma 8 and 9 therefore is in $O(n^3)$. Reductions corresponding to Lemma 10 and a node $v \in V_0$ can only become feasible when we remove incoming edges of v. This can happen at most n times for each node $v \in V_0$, before we remove it. The total run time for all reductions corresponding to Lemma 10 therefore is in $O(n^2)$. Reductions corresponding to Lemma 11 can only become feasible when all nodes of one subset have been removed. This can happen at most n times; hence any node will be moved to another subset at most n times. The total run time for all reductions corresponding to Lemma 11 therefore is in $O(n^2)$. □

We can now prove our main kernelization result.

Theorem 2. The part of the theorem for general instances follows directly from Lemma 7. The part for bipartite instances follows from Lemma 12 and Lemma 13 because the reduced bipartite parity game $G' = (V_0' \uplus V_1', E', p')$ satisfied $|V_0'| \leq 2^k \cdot \min\{k,p\}$ and $|V_1'| \leq k$. Since G' is bipartite, this implies that it contains at most $k2^k \cdot \min\{k,p\}$ edges. □

4 A Simple Exponential-Time Algorithm

A simple algorithm with run time $O(2^n)$ for the solution of parity games originates from the work of McNaughton [15] and was first presented for parity games by Zielonka [21]; see also Grädel et al. [9]. Algorithm $\mathbf{win}(G)$ receives a parity game G and returns the pair of winning sets $(\mathsf{win}_0(G) = W_0, \mathsf{win}_1(G) = W_1)$.

Algorithm $\mathbf{win}(G)$ is based on Lemmas 5 and 6. Let $p_{\max}$ be the largest priority in G and let $V_{p_{\max}}$ be the set of nodes with priority $p_{\max}$. Let $i = p_{\max} \pmod 2$ be the player who owns the highest priority. The algorithm first finds the winning sets (W_0', W_1') of the smaller game $G' = G - \mathsf{reach}_i(V_{p_{\max}})$ in a first recursive call. If $W_{\neg i}' = \emptyset$, then by Lemma 6 player i wins from all nodes of G and we are done. Otherwise, again by Lemma 6 we know that $W_{\neg i}' \subseteq \mathsf{win}_{\neg i}(G)$. The algorithm then finds the winning sets (W_0'', W_1'') of the smaller game $G'' = G - \mathsf{reach}_{\neg i}(W_{\neg i}')$ by a second recursive call. By Lemma 5, $\mathsf{win}_i(G) = W_i''$ and $\mathsf{win}_{\neg i}(G) = \mathsf{reach}_{\neg i}(W_{\neg i}') \cup W_{\neg i}'' = V(G) \setminus W_i''$.

Theorem 4. *Algorithm* $\mathbf{win}(G)$ *finds the winning sets of any parity game on* n *nodes in time* $O(2^n)$.

Proof. The correctness of the algorithm follows from Lemmas 5 and 6, as argued above. Let $T'(n)$ be the number of steps needed by algorithm $\mathbf{win}(G)$ to solve a game G on n nodes. Algorithm $\mathbf{win}(G)$ makes two

Algorithm 1 $\mathbf{win}(G)$

Input: A parity game $G = (V_0 \uplus V_1, E, p)$ with maximum priority $p_{\max}$.
Output: (W_0, W_1), where W_i is the winning set of player $i \in \{0, 1\}$.

```
1: if V = ∅ then
2:     return (∅, ∅)
3: i ← p_max (mod 2); j ← ¬i
4: (W'_0, W'_1) ← win(G − reach_i(V_{p_max}))
5: if W'_j = ∅ then
6:     (W_i, W_j) ← (V, ∅)
7: else
8:     (W''_0, W''_1) ← win(G − reach_j(W'_j))
9:     (W_i, W_j) ← (W''_i, V \ W''_i)
10: return (W_0, W_1)
```

recursive calls $\mathbf{win}(G')$ and $\mathbf{win}(G'')$ on games with at most $n-1$ nodes. Other than that, it performs only $O(n^2)$ operations. (The most time-consuming operations are the computations of the sets $\mathsf{reach}_i(V_{p_{\max}})$ and $\mathsf{reach}_j(W'_j)$.) Therefore, $T'(n) \leq 2T'(n-1) + O(n^2)$, which implies $T'(n) = O(2^n)$. □

5 Overview of the New Algorithms

Before we describe our new algorithms that lead to Theorems 1 and 3 in detail in Sect. 7 and Sect. 8, we present an overview of the main ideas. The algorithm **new-win**(G) by Jurdziński, Paterson, and Zwick [12] with run time $n^{O(\sqrt{n})}$ is a slight modification of the just described algorithm $\mathbf{win}(G)$. At the beginning of each recursive call it tests in time $O(n^\ell)$ if the parity game contains a dominion D of size at most $\ell = \lceil\sqrt{2n}\rceil$. If this is the case then D is removed and the remaining game is solved recursively. Else, the parity game is solved by the algorithm $\mathbf{win}(G)$, except that the recursive calls in lines 4 and 8 are made to **new-win**(G). Since this happens only when G does not contain a dominion of size at most ℓ, the dominion $\mathsf{reach}_j(W'_j)$ that is removed in line 8 has size greater than ℓ and hence, the second recursive call is to a substantially smaller game. Overall, this leads to the improved run time of $n^{O(\sqrt{n})}$.

Our new algorithms are based on a similar idea. Instead of simply searching for a dominion of size at most ℓ, our algorithm $\mathbf{new\text{-}win}_1(G)$ that leads to Theorem 1 searches for a dominion that contains at most $\ell = \lfloor\sqrt{2k}\rfloor$ nodes of the odd player, assuming without loss of generality that the odd player controls fewer nodes, i.e., $k = |V_1|$. If such a dominion is found then we remove it from the game and solve the remaining game recursively. Otherwise, we use the algorithm $\mathbf{win}(G)$ to solve the parity game, except that the recursive calls in lines 4 and 8 are made to $\mathbf{new\text{-}win}_1(G)$. It can happen that in the game to which the first recursive call in line 4 is made, the odd player controls again k nodes. We will show that in bipartite instances this cannot happen in two consecutive calls. For general instances we use that the observation that at least the number of different priorities decreases by one in the recursive call. Searching efficiently for a dominion that contains at most $\ell = \lfloor\sqrt{2k}\rfloor$ nodes of the odd player is more involved than simply searching for dominions whose total size is at most ℓ. We use multiple recursive calls of $\mathbf{new\text{-}win}_1$ to test if such a dominion exists, which makes the recursion of our algorithm and its analysis more complicated.

Our second algorithm leading to Theorem 3 is based on the same approach and inspired by the algorithm of Jurdziński, Paterson, and Zwick [12]. In this case we let s_j, for some $j \in \mathbb{N}$, equal the number of nodes with out-degree at most j. We separate the nodes into s_j nodes with out-degree at most j and $n - s_j$ nodes with out-degree larger than j and, at the beginning of each iteration, search for and remove dominions that contain at most $\ell = \lceil\sqrt{2(n-s_j)}\rceil$ nodes with out-degree larger than j and at most $s = \lceil\sqrt{s_j \cdot \log_j s_j}\rceil$ nodes with out-degree at most j. This algorithm runs in time $n^{O\left(\sqrt{n-s_j}+\sqrt{\frac{s_j}{\log_j s_j}}\right)}$, which implies Theorem 3. □

6 Finding Small Dominions

We now describe how dominions with the previously discussed properties can be found. Let $G = (V_0 \uplus V_1, E, p)$ be a parity game. Recall that for $i \in \{0,1\}$, a set $D \subseteq V$ is an i-dominion if player i can win from every node of D without ever leaving D, regardless of the strategy of player $\neg i$. Note that any i-dominion must be i-closed. A set $D \subseteq V$ is a *dominion* if it is either a 0-dominion or a 1-dominion. By prefix independence of parity games, the winning set $\mathsf{win}_i(G)$ of player i is an i-dominion.

For $k, p \in \mathbb{N}$, let $T(k)$ denote the maximum number of steps needed to solve a bipartite parity game $G = (V_0 \uplus V_1, E, p)$ and let $T(k,p)$ denote the maximum number of steps needed to solve a general parity game $G = (V_0 \uplus V_1, E, p)$ with $|V_1| = k$ and $p = |\{p(v) \mid v \in V\}|$ using some fixed algorithm. For $k, p, \ell \in \mathbb{N}$, let $\mathsf{dom}_k(\ell)$ denote the maximum number of steps required to find a dominion D with $|V_1 \cap D| \leq \ell$ in a bipartite parity game $G = (V_0 \uplus V_1, E, p)$ with $|V_1| = k$ and let $\mathsf{dom}_{k,p}(\ell)$ denote the maximum number of steps required to find a dominion D with $|V_1 \cap D| \leq \ell$ in a general parity game $G = (V_0 \uplus V_1, E, p)$ with $|V_1| = k$ and $p = |\{p(v) \mid v \in V\}|$, or to determine that no such dominion exists.

We will in the analysis of run times make the assumption that computation and removal of reachability sets as well as kernelization are elementary operation and can therefore be performed in time $O(1)$. To obtain the actual run times of our algorithms we will in the end multiply the computed run times by a factor corresponding to the time needed for these operations.

Lemma 14. *For $k \geq 4$, $\mathsf{dom}_k(\ell) = O(k^\ell \cdot T(\ell))$ and $\mathsf{dom}_{k,p}(\ell) = O(k^\ell \cdot T(\ell, p))$.*

Proof. There are $O(k^\ell)$ sets $V_D \subseteq V_1$ with $|V_D| = \ell$. We argue below that for each such set V_D, one can determine whether or not there exists a dominion D with $D \cap V_1 \subseteq V_D$ by solving two parity games that are sub-games of G, i.e., these games arise from G by removing some of the nodes. This implies the lemma because each of these sub-games can be solved in time $T(\ell)$ or $T(\ell, p)$ for bipartite or general parity games, respectively.

Let $V_D \subseteq V_1$ be a set with $|V_D| = \ell$. We will now show how to check if there exists an i-dominion D with $D \cap V_1 \subseteq V_D$. If such an i-dominion D exists, then it is i-closed. Therefore, it does not contain any node $v \in V$ from which player $\neg i$ can reach a node in $V_1 \setminus V_D$. Let $V' = V(G) \setminus \mathsf{reach}_{\neg i}(V_1 \setminus V_D)$ be the set of nodes from which player $\neg i$ cannot force to reach a node in $V_1 \setminus V_D$; the set V' can therefore be computed by computing and removing a reachability set, which as we assumed is an elementary operations. We then have $D \subseteq V'$, and since no node in $V_1 \setminus V_D$ can be part of V', it holds that $V' \cap V_1 \subseteq V_D$. Since V' is an i-closed set, the game $G - \mathsf{reach}_{\neg i}(V_1 \setminus V_D)$ is well defined. Let $\mathsf{win}_i(V')$ be the winning set of player 1 in the game $G - \mathsf{reach}_{\neg i}(V_1 \setminus V_D)$. Then $\mathsf{win}_i(V')$ is an i-dominion that contains D.

This shows that for each set $V_D \subseteq V_1$ with $|V_D| = \ell$ we only need to compute for $i \in \{0,1\}$ the sets $V_i' = V \setminus \mathsf{reach}_{\neg i}(V_1 \setminus V_D)$ of nodes from which player $\neg i$ cannot force to enter $V_1 \setminus V_D$ and compute the winning sets of the game $G - \mathsf{reach}_{\neg i}(V_1 \setminus V_D)$ to determine whether or not there exists a dominion D with $D \cap V_1 \subseteq V_D$. □

With the algorithm described in Lemma 14 we can find a dominion D such that $|D \cap V_1| \leq \ell$ if such a dominion exists. We denote this algorithm by $\mathbf{dominion}_1(G, \ell)$ and assume that it returns either the pair (D, i) if an i-dominion D is found, or $(\emptyset, -1)$ if not.

We will give the pseudocode for algorithm $\mathbf{dominion}(G, \ell, s)$. In the pseudocode, let U_m denote the set of marked nodes from U and let $\mathbf{king}(U, \mathrm{strategy}_i)$ denote an execution of the algorithm by King et al. [13] that determines the winners of the sub-game G restricted to U with a given strategy for player i.

7 New Algorithms for Solving Parity Games

We present the algorithm $\mathbf{new\text{-}win}_1(G)$ discussed in Sect. 5 in detail. Let $G = (V_0 \uplus V_1, E, p)$ with $|V_1| = k$ be a parity game with p distinct priorities.

The algorithm $\mathbf{new\text{-}win}_1$ starts by trying to find a "small" dominion D, where small means $|D \cap V_1| \leq \ell$, where $\ell = \lfloor \sqrt{2k} \rfloor$ is a parameter chosen to minimize the run time of the algorithm. If such an i-dominion is

Algorithm 2 dominion(G, ℓ, s)

Input: A parity game $G = (V_0 \uplus V_1, E, p)$ and $\ell, s \in \{0, \ldots, |V(G)|\}$.
Output: An i-dominion (D, i) for $i \in \{0, 1\}$ or $(\emptyset, -1)$ if no dominion is found.

Fix a total order $\prec$ on the nodes of G.
For each $u \in V(G)$ sort the edges emanating from u by $\prec$ on their respective endpoint.
for $i \in \{0, 1\}$ **do**
  for $v \in V_i$ **do**
    for $\langle a_1, \ldots, a_\ell, b_1, \ldots, b_s \rangle \in \{1, \ldots, |V(G)|\}^\ell \times \{1, \ldots, j\}^s$ **do**
      $r_1 = 1, r_2 = 1, U = \{v\}, U_m = \emptyset$
      while $U \neq \emptyset$, $r_1 \leq \ell$ and $r_2 \leq s$ **do**
        Choose $u = \min(U, \prec)$.
        $U = U \setminus \{u\}$, $U_m = U_m \cup \{u\}$.
        if $u \in V_i$ **then**
          if $|\delta^+(u)| > j$ **then**
            if $|\delta^+(u)| \leq r_1$ **then**
              Let $e = (u, w)$ be the a_{r_1}-th outgoing edge of u.
              $U = U \cup (\{w\} \setminus U_m)$, $r_1 = r_1 + 1$, $\text{strategy}_i(u) = w$.
            else
              $U = \emptyset$, $U_m = \emptyset$.
          else
            if $|\delta^+(u)| \leq r_2$ **then**
              Let $e = (u, w)$ be the b_{r_2}-th outgoing edge of u.
              $U = U \cup (\{w\} \setminus U_m)$, $r_2 = r_2 + 1$, $\text{strategy}_i(u) = w$.
            else
              $U = \emptyset$, $U_m = \emptyset$.
        else
          $U = U \cup (N^+(u) \setminus U_m)$
      if $U_m \neq \emptyset$ contains at most ℓ high out-degree nodes and at most s low out-degree nodes **then**
        $(W_0, W_1) =$ **king**$(U, \text{strategy}_i)$.
        if $W_i = U$ **then**
          return (U, i)
return $(\emptyset, -1)$.

found, then we remove it together with its i-reachability set from the game and solve the remaining game recursively. If no small dominion is found, then **new-win**$_1$ simply calls algorithm **old-win**$_1$, which is almost identical to algorithm **win**. The only difference between **old-win**$_1$ and **win** is that its recursive calls are made to **new-win**$_1$ and not to itself.

The recursion stops once the number of odd nodes is at most 4, in which case we will test each of the at most $((p+1)^4)^4$ (due to the size of our kernel) different strategies for player 1 in constant time. We will call this brute force method **solve**(G). We will also kernelize using the reduction rules described in Sect. 3. We will call the kernelization subroutine **kernel**(G). The pseudocode of **new-win**$_1(G)$ can be found in Sect. 9.

The correctness of the algorithm follows analogously to the correctness of **win**(G). We analyze the run time of **new-win**$_1(G)$ and prove Theorem 1 in section 10.

8 Out-degree based Algorithm

We now describe our second algorithm **new-win**$_2(G, j)$. In order to describe it, let $j \in \mathbb{N}$ and let s_j denote the number of nodes of out-degree at most j. **new-win**$_2(G, j)$ is then almost identical to **new-win**$_1(G)$, but instead of dominions that contains at most $\ell' = \lfloor \sqrt{2k} \rfloor$ nodes of the odd player, we search for and delete

dominions that contain at most $\ell = \lceil\sqrt{2(n-s_j)}\rceil$ nodes with out-degree larger than j and at most $s = \lceil\sqrt{s_j \cdot \log_j s_j}\rceil$ nodes with out-degree at most j. This algorithm has a run time of $n^{O\left(\sqrt{n-s_j}+\sqrt{\frac{s_j}{\log_j s_j}}\right)}$, which implies Theorem 3.

In the following let us assume $j = \arg\min_{1\leq j'\leq n}\left\{\sqrt{n-s_{j'}} + \sqrt{\frac{s_{j'}}{\log_{j'} s_{j'}}}\right\}$. We say that a node v has *high out-degree* if $|\delta^+(v)| > j$ and *low out-degree* otherwise. For $n, z, \ell, s \in \mathbb{N}$, let $\mathsf{dom}_{n,z}(\ell, s)$ denote the maximum number of steps required to find a dominion D with at most ℓ high out-degree and at most s low out-degree nodes in a parity game $G = (V_0 \uplus V_1, E, p)$ with n nodes out of which z are high out-degree nodes, or to determine that no such dominion exists.

Lemma 15. *For all values of $\ell, s \in \{0, \ldots, n\}$, it holds*

$$\mathsf{dom}_{n,z}(\ell, s) = O\left(n^{\ell+1} j^s (\ell+s)^2 \cdot \max\{1, \log(\ell+s)\}\right) = O\left(n^{\ell+4} j^s\right).$$

Proof. Fix an arbitrary total order $\prec$ on $V(G)$. Let $u \in V(G)$ be a node of G and let $(u, v_1), \ldots, (u, v_{|\delta^+(u)|})$ be the edges emanating from u, where $v_i \prec v_{i+1}$ for all $i \in \{1, \ldots |\delta^+(v)|-1\}$; we call (u, v_i) the *i-th outgoing edge of u*. The algorithm generates at most $O(n \cdot n^\ell j^s)$ 0-closed sets of nodes that contain at most ℓ nodes with an out-degree greater than j and at most s nodes with an out-degree at most j, which are candidates for being 0-dominions. For every node $v \in V$ and every sequence $\langle a_1, \ldots, a_\ell, b_1, \ldots, b_s\rangle \in \{1, \ldots, n\}^\ell \times \{1, \ldots, j\}^s$ construct a set $U \subseteq V$ as follows. Start with $U = \{v\}$ and $r_1 = 1, r_2 = 1$. Nodes added to U are initially unmarked. As long as there is still an unmarked node in U, pick the smallest such node $u \in U$ with respect to $\prec$ and mark it.

- If $u \in V_0$ and u has high out-degree then add the endpoint of the a_{r_1}-th outgoing edge of u to U (if it is not already present in U) and increment r_1.
- If $u \in V_0$ and u has low out-degree then add the endpoint of the b_{r_2}-th outgoing edge of u to U (if it is not already present in U) and increment r_2.
- If $u \in V_1$ then add the endpoints of all outgoing edges of u that are not yet part of U to U.

If at some stage U contains either more than ℓ nodes with high out-degree or more than s nodes with low out-degree, or the endpoint of the i-th outgoing edge of some node v with out-degree $|\delta^+(v)| < i$ should be added to U, then discard the set U and restart the construction with the next sequence. If the process above ends without discarding U, then a 0-closed set containing at most ℓ high out-degree and at most s low out-degree nodes has been found. Furthermore, for every node $u \in U \cap V_0$, one of the outgoing edges of u was selected. This corresponds to a suggested strategy for player 0 in the game G restricted to the set U.

Our algorithm therefore considers by exhaustive search all 0-closed sets containing at most ℓ high out-degree and at most s low out-degree nodes, and for each set considers all possible positional strategies for player 0. Using an algorithm of King et al. [13] we can check in time $O((\ell+s)^2 \log(\ell+s))$ time whether a given pair of set U and proposed strategy is indeed a winning strategy for player 0 from all nodes of U. Thus, if there is a 0-dominion containing at most ℓ high out-degree and at most s low out-degree nodes, then the algorithm will find one. Finding 1-dominions can be done in an analogous manner. □

With the just described algorithm, we can find a dominion D with at most ℓ nodes with high out-degree and at most s nodes with low out-degree if such a dominion exists. We denote this algorithm by $\mathbf{dominion}_2(G, \ell, s)$, and suppose that it returns either the pair (D, i) if such an i-dominion D is found, or $(\emptyset, -1)$ if not.

The algorithm $\mathbf{new\text{-}win}_2$ starts by trying to find a "small" dominion D, where small means that D contains at most ℓ nodes with out-degree greater than j and at most s nodes with out-degree at most j, where $\ell = \lceil\sqrt{2(n-s_j)}\rceil$ and $s = \lceil\sqrt{s_j \cdot \log_j s_j}\rceil$ are parameters chosen to minimize the run time of the whole algorithm. If such an i-dominion is found, then we remove it together with its i-reachability set from the game and solve the remaining game recursively. If no small dominion is found, then $\mathbf{new\text{-}win}_2$ simply

calls algorithm **old-win**$_2$, which is almost identical to algorithm **win**. The only difference between **old-win**$_2$ and **win** is that its recursive calls are made to **new-win**$_2$ and not to itself.

The recursion stops once the number of nodes with out-degree at most j and the number of nodes with out-degree greater than j are both at most 3, in which case we will test all of the at most constant different strategies for the two players in constant time. We will call this brute force method **solve(**G**)**. The pseudocode of **new-win**$_2$ can be found in Sect. 9.

The correctness of **new-win**$_2$ follows analogously to the correctness of the simple algorithm **win**. We analyze the run time of **new-win**$_2$ and prove Theorem 3 in Sect. 10.

We can now proove Corollary 2.

Corollary 2. First consider a parity game on n nodes played on a graph with maximum out-degree Δ. Then $s_\Delta = n$ and

$$\sqrt{n - s_\Delta} + \sqrt{\frac{s_\Delta}{\log_\Delta s_\Delta}} = \sqrt{\frac{n}{\log_\Delta n}} = \sqrt{\frac{\log(\Delta)n}{\log n}} \ .$$

Now the first part of the corollary follows immediately from Theorem 3.

Let us now consider the case that the average out-degree is Δ and let $z = \log(n)\Delta$. Then Markov's inequality implies $s_z \geq (1 - 1/\log(n))n$. Hence,

$$\sqrt{n - s_z} + \sqrt{\frac{s_z}{\log_z s_z}} \leq \sqrt{\frac{n}{\log(n)}} + \sqrt{\frac{n}{\log_z n}} = \sqrt{\frac{n}{\log(n)}} + \sqrt{\frac{\log(\log(n)\Delta)n}{\log n}} \ .$$

Now the second part of the corollary follows immediately from Theorem 3. □

9 Pseudocode for Algorithms new-win

We will now give the pseudocode for the algorithms **new-win**$_1(G)$ and **new-win**$_2(G, j)$ together with their subroutines **old-win**$_1(G)$ and **old-win**$_2(G, j)$. In the pseudocodes, we call a function **solve**(G). This function denotes a bruteforce method to solve parity games and is only used on very small games.

Algorithm 3 new-win$_1(G)$

Input: A parity game $G = (V_0 \uplus V_1, E, p)$.
Output: A partition (W_0, W_1) of V, where W_i is the winning set of player $i \in \{0, 1\}$.

1: $k \leftarrow |V_1|; \ell \leftarrow \left\lfloor \sqrt{2k} \right\rfloor$; $G = \textbf{kernel}(G)$
2: **if** $k \leq 4$ **then return solve**(G)
3: $(D, i) \leftarrow \textbf{dominion}_1(G, \ell)$
4: **if** $D = \emptyset$ **then**
5: $\quad (W_0, W_1) \leftarrow \textbf{old-win}_1(G)$
6: **else**
7: $\quad (W'_0, W'_1) \leftarrow \textbf{new-win}_1(G - \mathsf{reach}_i(D))$
8: $\quad (W_{\neg i}, W_i) \leftarrow (W'_{\neg i}, V \setminus W'_{\neg i})$
9: **return** (W_0, W_1)

Algorithm 4 old-win$_1(G)$

Input: A parity game $G = (V_0 \uplus V_1, E, p)$.
Output: A partition (W_0, W_1) of V, where W_i is the winning set of player $i \in \{0, 1\}$.

$G = \textbf{kernel}(G)$
$i \leftarrow p_{\max} \pmod 2$
$(W'_0, W'_1) \leftarrow \textbf{new-win}_1(G - \mathsf{reach}_i(V_{p_{\max}}))$
if $W'_{\neg i} = \emptyset$ **then**
    $(W_i, W_{\neg i}) \leftarrow (V, \emptyset)$
else
    $(W''_0, W''_1) \leftarrow \textbf{new-win}_1(G - \mathsf{reach}_{\neg i}(W'_{\neg i}))$
    $(W_i, W_{\neg i}) \leftarrow (W''_i, V \setminus W''_i)$
return (W_0, W_1)

Algorithm 5 new-win$_2(G, j)$

Input: A parity game $G = (V_0 \uplus V_1, E, p)$ and $j \in \{1, \dots |V|\}$.
Output: A partition (W_0, W_1) of V, where W_i is the winning set of player $i \in \{0, 1\}$.

$s_j \leftarrow |\{v \in V || \delta^+(v)| \leq j\}|; \ell \leftarrow \left\lceil \sqrt{2(n - s_j)} \right\rceil ; s \leftarrow \lceil \sqrt{s_j \cdot \log_j s_j} \rceil$
if $s_j \leq 3$ and $n - s_j \leq 3$ **then return solve**(G)
$(D, i) \leftarrow \textbf{dominion}_2(G, \ell, s)$
if $D = \emptyset$ **then**
    $(W_0, W_1) \leftarrow \textbf{old-win}_2(G, j)$
else
    $(W'_0, W'_1) \leftarrow \textbf{new-win}_2(G - \mathsf{reach}_i(D), j)$
    $(W_{\neg i}, W_i) \leftarrow (W'_{\neg i}, V \setminus W'_{\neg i})$
return (W_0, W_1)

Algorithm 6 old-win$_2(G, j)$

Input: A parity game $G = (V_0 \uplus V_1, E, p)$.
Output: A partition (W_0, W_1) of V, where W_i is the winning set of player $i \in \{0, 1\}$.

$i \leftarrow p_{\max} \pmod 2$
$(W'_0, W'_1) \leftarrow \textbf{new-win}_2(G - \mathsf{reach}_i(V_{p_{\max}}), j)$
if $W'_{\neg i} = \emptyset$ **then**
    $(W_i, W_{\neg i}) \leftarrow (V, \emptyset)$
else
    $(W''_0, W''_1) \leftarrow \textbf{new-win}_2(G - \mathsf{reach}_{\neg i}(W'_{\neg i}), j)$
    $(W_i, W_{\neg i}) \leftarrow (W''_i, V \setminus W''_i)$
return (W_0, W_1)

10 Analysis of the Run Time

We will show that algorithm **new-win**(G) has a run time of $O(p \cdot m \cdot n) \cdot (p + k)^{O(\sqrt{k})}$ on general instances and in time $O(n^3) \cdot k^{O(\sqrt{k})}$ on bipartite instances. We will also show that algorithm **new-win**(G, j) has a run time of $n^{O\left(\sqrt{n - s_j} + \sqrt{\frac{s_j}{\log_j s_j}}\right)}$, where s_j the number of nodes in G with out-degree at most j.

Note that the part $O(pmn)$ of the run time comes from the reduction of the instance and the computation and removal of reachability sets of found dominions. Since we do both of these often, we assume them to be elementary computations with computation time $O(1)$ and show that the total run time remaining is $(p+k)^{O(\sqrt{k})}$. In bipartite instances we need $O(n^3)$ time to reduce the instance and to compute and remove reachability sets. We will show that the run time on bipartite instances is $k^{O(\sqrt{k})}$, when computation and removal of reachability sets and the reductions are viewed as an elementary operation. Let $T(k,p)$ denote the time required by algorithm **new-win** on a game $G = (V_0 \uplus V_1, E, p)$ with $|V_1| = k$ and $p = |\{p(v) \mid v \in V\}|$, when reduction of an instance and computation and removal of reachability sets of found dominions are viewed as elementary computations and have run time $O(1)$.

Lemma 16. *The following recurrence relation holds:*

$$(a) \quad T(k,p) \leq \max\{T(k-1,p), T(k,p-1) + T(k-\ell,p)\} + \mathsf{dom}_{k,p}(\ell) + O(1) \ .$$

Proof. Algorithm **new-win**(G) tries to find dominions D with $|D \cap V_1| \leq \ell = \lfloor\sqrt{2k}\rfloor$. By definition this takes at most $\mathsf{dom}_{k,p}(\ell)$ time on general instances. If a (non-empty) dominion is found, then the algorithm simply proceeds on the remaining game, which has at most $k-1$ odd nodes, and thus it solves this game in time bounded by $T(k-1,p)$. Otherwise, a call to **old-win**(G) is made. This results in a call to **new-win**$(G - \mathsf{reach}_i(V_{p_{\max}}))$. In this case the call takes at most $T(k,p-1)$ time because we removed all nodes with the highest priority. If the set W'_j returned by the call is empty, then we are done. Otherwise, $W'_j = \mathsf{win}_j(G - \mathsf{reach}_i(V_{p_{\max}}))$ and $W'_j \subseteq \mathsf{win}_j(G)$ by Lemma 4. Therefore, W'_j is a j-dominion of G. We are in the case that there is no dominion D with $|D \cap V_1| \leq \ell$ in G. Thus, $|W'_j \cap V_1| > \ell$, and hence the second recursive call **new-win**$(G - \mathsf{reach}_j(W'_j))$ takes time at most $T(k - \lceil\ell\rceil, p)$. Consequently,

$$T(k,p) \leq \max\{T(k-1,p), T(k,p-1) + T(k-\lceil\ell\rceil,p)\} + \mathsf{dom}_{k,p}(\ell) + O(1). \qquad \square$$

Let $T(k,\mathsf{even})$ and $T(k,\mathsf{odd})$ denote the time required by algorithm **new-win** on a bipartite game $G = (V_0 \uplus V_1, E, p)$ with $|V_1| = k$ when the largest priority is even respectively odd, when computation and removal of reachability sets of found dominions and the reductions are viewed as elementary computations. We denote by $T(k)$ the time required by algorithm **new-win** on any bipartite game with $|V_1| = k$; thus $T(k) \leq \max\{T(k,\mathsf{even}), T(k,\mathsf{odd})\}$.

Lemma 17. *The following recurrence relations hold:*

$$\begin{aligned}
(b_1) \quad & T(k,\mathsf{odd}) \leq T(k-1) + T(k-\ell) + \mathsf{dom}_k(\ell) + O(1), \\
(b_2) \quad & T(k,\mathsf{even}) \leq \max\{T(k,\mathsf{odd}), T(k-1)\} + T(k-\ell) + \mathsf{dom}_k(\ell) + O(1) \\
& \qquad\qquad \leq T(k-1) + 2T(k-\ell) + 2\mathsf{dom}_k(\ell) + O(1), \\
(b_3) \quad & T(k) \leq T(k-1) + 2T(k-\ell) + 2\mathsf{dom}_k(\ell) + O(1) \ .
\end{aligned}$$

Proof. From the definition it follows directly that $T(k) \leq \max\{T(k,\mathsf{even}), T(k,\mathsf{odd})\}$. Showing (b_1) and (b_2) therefore yields (b_3). Algorithm **new-win**(G) tries to find dominions D with $|D \cap V_1| \leq \ell = \lceil\sqrt{2k}\rceil$. By definition this takes at most $\mathsf{dom}_k(\ell)$ time on bipartite instances. If a (non-empty) dominion is found, then the algorithm simply proceeds on the remaining game, which has at most $k-1$ odd nodes, and the remaining run time is therefore at most $T(k-1)$. Otherwise, a call to **old-win**(G) is made. This results in a call to **new-win**$(G - \mathsf{reach}_i(V_{p_{\max}}))$. Here we have to distinguish whether the highest priority is odd or even.

If the highest priority is odd then, by Lemma 10, the set $\mathsf{reach}_1(V_{p_{\max}}) \cap V_1$ is non-empty and the call takes at most $T(k-1)$ time.

In case the highest priority is even, we either have $\mathsf{reach}_0(V_{p_{\max}}) \cap V_1 \neq \emptyset$ or $\mathsf{reach}_0(V_{p_{\max}}) = V_{p_{\max}}$ in which case Lemma 11 yields that in $G - \mathsf{reach}_i(V_{p_{\max}})$ the highest priority has to be odd. Therefore, this call needs time at most $\max\{T(k,\mathsf{odd}), T(k-1)\}$.

If the set W'_j returned by the call is empty, then we are done. Otherwise, $W'_j = \mathsf{win}_j(G - \mathsf{reach}_i(V_{p_{\max}}))$ and this is part of $\mathsf{win}_j(G)$ by Lemma 4. Therefore, W'_j is a j-dominion of G. We are in the case that there

is no dominion D with $|D \cap V_1|$ at most ℓ in G, so we know that $|W'_j \cap V_1| > \ell$, and therefore the second recursive call **new-win**$(G - \mathsf{reach}_j(W'_j))$ takes at most $T(k-\ell)$ time. Thus, we obtain

$$T(k, \mathsf{odd}) \leq T(k-1) + T(k-\ell) + \mathsf{dom}_k(\ell)$$

and

$$\begin{aligned} T(k, \mathsf{even}) &\leq \max\{T(k, \mathsf{odd}), T(k-1)\} + T(k-\ell) + \mathsf{dom}_k(\ell) \\ &\leq T(k-1) + 2T(k-\ell) + 2\mathsf{dom}_k(\ell), \end{aligned}$$

which yields $T(k) \leq T(k-1) + 2T(k-\ell) + 2\mathsf{dom}_k(\ell)$. □

For $j \in \mathbb{N}_0$, let $T'(s_j, n-s_j)$ denote the time required by algorithm **new-win** on a game G on n nodes of which s_j nodes have out-degree at most j.

Lemma 18. *The following recurrence relation holds:*

$$\begin{aligned} (c) \quad T'(s_j, n-s_j) \leq{}& \max\{T'(s_j-1, n-s_j), T'(s_j, n-s_j-1)\} \\ &+ \max\{T'(s_j-s, n-s_j), T'(s_j, n-s_j-\ell)\} + \mathsf{dom}_{n,n-s_j}(\ell, s) + O(1) \ . \end{aligned}$$

Proof. Algorithm **new-win**(G, j) tries to find dominions D containing at most s nodes with out-degree at most j and at most ℓ nodes with out-degree greater than j. By definition this takes at most $\mathsf{dom}_{n,n-s_j}(\ell, s)$ time. If a (non-empty) dominion is found, then the algorithm simply proceeds on the remaining game, which has at most $n-1$ nodes, and the remaining time is therefore at most $\max\{T'(s_j-1, n-s_j), T'(s_j, n-s_j-1)\}$. Otherwise, a call to **old-win**(G, j) is made. This results in a call to **new-win**$(G - \mathsf{reach}_i(V_{p_{\max}}), j)$, this call is to a game with fewer nodes and can be solved in time bounded by

$$\max\{T'(s_j-1, n-s_j), T'(s_j, n-s_j-1)\} \ .$$

If the set W'_k returned by the call is empty, then we are done. Otherwise, $W'_k = \mathsf{win}_k(G - \mathsf{reach}_i(V_{p_{\max}}))$, and $W'_k \subseteq \mathsf{win}_k(G)$ by Lemma 4. Therefore, W'_k is a k-dominion of G. We are in the case that there is no dominion D containing at most s nodes with out-degree at most j and at most ℓ nodes with out-degree greater than j, so W'_k either contains more than s nodes with out-degree at most j or more than ℓ nodes with out-degree greater than j, and therefore the second recursive call **new-win**$(G - \mathsf{reach}_k(W'_k))$ takes time bounded by $\max\{T'(s_j-s, n-s_j), T'(s_j, n-s_j-\ell)\}$.

All other computations can be done in constant time. Thus, we obtain

$$\begin{aligned} T'(s_j, n-s_j) \leq{}& \max\{T'(s_j-1, n-s_j), T'(s_j, n-s_j-1)\} \\ &+ \max\{T'(s_j-s, n-s_j), T'(s_j, n-s_j-\ell)\} + \mathsf{dom}_{n,n-s_j}(\ell, s) + O(1) \ . \end{aligned}$$ □

We analyze recurrences (a) and (b) with $\ell = \lfloor\sqrt{2k}\rfloor$ in Theorem 5 in Sect. 11, which eventually shows that $T(k,p) \leq (p+k)^{O(\sqrt{k})}$ and $T(k) \leq k^{O(\sqrt{k})}$, and recurrence (c) with $\ell = \left\lceil \sqrt{2(n-s_j)} \right\rceil$ and $s = \left\lceil \sqrt{\frac{s_j}{\log_j s_j}} \right\rceil$ in Theorem 22 in Sect. 11, which eventually shows that $T'(s_j, n-s_j) \leq n^{O\left(\sqrt{n-s_j} + \sqrt{\frac{s_j}{\log_j s_j}}\right)}$. This completes the analysis of the run time of **new-win**(G) and **new-win**(G, j), and it proves Theorem 1 and Theorem 3. □

11 Recurrence Relation Computations

In this section we analyze the recurrence relations used to bound the run time of **new-win**.

Theorem 5. *For $k \in \mathbb{N}$ and $\ell = \lfloor\sqrt{2k}\rfloor$, we obtain*

$$\begin{aligned} T(k,p) &= (p+k)^{O(\sqrt{k})}, \\ T(k) &= k^{O(\sqrt{k})}. \end{aligned}$$

To prove Theorem 5, we first establish some lemmas.

Lemma 19. *For $k, p \in \mathbb{N}$ and $\ell = \lfloor\sqrt{2k}\rfloor$, it holds*

$$\begin{aligned} T(k,p) &\leq 2(k+p)^{\lfloor\sqrt{2k}\rfloor} \cdot \mathsf{dom}_{k,p}(\left\lfloor\sqrt{2k}\right\rfloor) \quad \textit{and} \\ T(k) &\leq 2(2k)\left\lfloor\sqrt{2k}\right\rfloor \cdot \mathsf{dom}_k(\left\lceil\sqrt{2k}\right\rceil) \ . \end{aligned}$$

Proof. For every pair of integers k and p we construct binary trees $T_{k,p}$ and T_k in the following way. The root of $T_{k,p}$ is labeled by k and $k+p$ and the root of T_k is labeled by k. A node labeled by a number $k > 4$ has two children: in $T_{k,p}$ a left child labeled by k and $k+p-1$ and a right child labeled by $k - \left\lceil\sqrt{2k}\right\rceil$ and $p+k-\left\lceil\sqrt{2k}\right\rceil$. In T_k a left child labeled by k' and a right child labeled by $k-\left\lceil\sqrt{2k}\right\rceil$. A node labeled by k' in T_k has two children: a left child labeled by $k-1$ and a right child labeled by $k-\left\lceil\sqrt{2k}\right\rceil$. Nodes labeled by a number $k \leq 4$ are leaves. A node labeled by k and $k+p$ has a cost of $\mathsf{dom}_{k,p}(\left\lfloor\sqrt{2k}\right\rfloor)$ associated with it and a node labeled by k or k' has a cost of $\mathsf{dom}_k(\left\lceil\sqrt{2k}\right\rceil))$ associated with it. It follows from Lemma 16 and Lemma 17 that the sum of the costs of the nodes of $T_{k,p}$ and T_k, is an upper bound on $T(k,p)$ and $T(k)$, respectively. Clearly, the length of every path from the root to a leaf is at most $p+k+1$ in $T_{k,p}$ and $2k$ in T_k. We say that such a path makes a right turn when it descends from a node to its right child. We next claim that each such path makes at most $\left\lfloor\sqrt{2k}\right\rfloor$ right turns. This follows immediately from the observation that the function $f(n) = n - \lceil\sqrt{2n}\rceil$ can be iterated on $2k$ at most $\left\lfloor\sqrt{2k}\right\rfloor$ times before reaching the value of 4 or less. This observation can be proved by induction, based on the fact that if $\frac{1}{2}j^2 < n \leq \frac{1}{2}(j+1)^2$, then $n - \lceil\sqrt{2n}\rceil \leq \frac{1}{2}j^2$. (Initially we have $j = \left\lfloor\sqrt{2k}\right\rfloor$ and finally, with $1 \leq n \leq 4$, we have $j \geq 1$.) As each leaf of $T_{k,p}$ and T_k is determined by the positions of the right turns on the path leading to it from the root, we get that the number of leaves is at most $\binom{k+p}{\lfloor\sqrt{2k}\rfloor}$ in $T_{k,p}$ and at most $\binom{2k}{\lfloor\sqrt{2k}\rfloor}$ in T_k. The total number of nodes is therefore at most at most $2\binom{k+p}{\lfloor\sqrt{2k}\rfloor} \leq 2(k+p)\left\lfloor\sqrt{2k}\right\rfloor$ in $T_{k,p}$ and at most $2\binom{2k}{\lfloor\sqrt{2k}\rfloor} \leq 2(2k)^{\lfloor\sqrt{2k}\rfloor}$ in T_k. As the cost of each node is at most $\mathsf{dom}_{k,p}(\left\lfloor\sqrt{2k}\right\rfloor)$ in $T_{k,p}$ and at most $\mathsf{dom}_k(\left\lceil\sqrt{2k}\right\rceil)$ in T_k, we immediately get

$$\begin{aligned} T(k,p) &\leq 2(k+p)^{\lfloor\sqrt{2k}\rfloor} \cdot \mathsf{dom}_{k,p}(\left\lfloor\sqrt{2k}\right\rfloor) \quad \text{and,} \\ T(k) &\leq 2(2k)^{\lfloor\sqrt{2k}\rfloor} \cdot \mathsf{dom}_k(\left\lceil\sqrt{2k}\right\rceil) \ . \end{aligned}$$

□

We obtain together with Lemma 14 that

$$T(k,p) \leq 2(k+p)^{\lfloor\sqrt{2k}\rfloor} \cdot O\Big(k^{\lfloor\sqrt{2k}\rfloor} \cdot T(\left\lfloor\sqrt{2k}\right\rfloor, p)\Big),$$

as well as

$$T(k) \leq 2(2k)^{\left\lfloor\sqrt{2k}\right\rfloor} \cdot O\left(k^{\lceil\sqrt{2k}\rceil} T(\left\lceil\sqrt{2k}\right\rceil)\right).$$

Lemma 20. *Suppose that*

$$T(k,p) \leq 2(k+p)^{\lfloor\sqrt{2k}\rfloor} \cdot O\left(k^{\lfloor\sqrt{2k}\rfloor} \cdot T(\left\lfloor\sqrt{2k}\right\rfloor, p)\right)$$

and that $T(\ell,q) \leq c' \cdot (q+\ell)^{8\lfloor\sqrt{2\ell}\rfloor}$ *for some constant* $c' \in \mathbb{R}$ *and for all pairs* $(\ell,q) \in \{1,\dots,4\} \times \mathbb{N}$. *Then there exist constants* $c_1 \geq c', c_2 \geq 8$ *such that for all* $k \in \mathbb{N}$,

$$T(k,p) \leq c_1 \cdot (p+k)^{c_2\lfloor\sqrt{2k}\rfloor} \ .$$

Proof. Since we have $T(k,p) \leq 2(k+p)^{\lfloor\sqrt{2k}\rfloor} \cdot O\left(k^{\lfloor\sqrt{2k}\rfloor} \cdot T(\left\lfloor\sqrt{2k}\right\rfloor, p)\right)$, there exists a constant $c_1' > 0$ such that $T(k,p) \leq 2(k+p)^{\lfloor\sqrt{2k}\rfloor} \cdot c_1' k^{\lfloor\sqrt{2k}\rfloor} \cdot T(\left\lfloor\sqrt{2k}\right\rfloor, p)$. Let $\alpha_k = \frac{\lfloor\sqrt{2k}\rfloor}{\left\lfloor\sqrt{2\lfloor\sqrt{2k}\rfloor}\right\rfloor}$. Then for $k \geq 5$, it holds $\alpha_k \geq 1.5 > 1$. Let $c_1 = \max\{c_1', c'\}$, and let $c_2 = 6 + 3\log(2c_1)$. Suppose, for sake of contradiction, that the statement of the lemma does not hold for this choice of (c_1, c_2). Then there exists a pair $(k',p') \in \mathbb{N} \times \mathbb{N}$ for which $T(k',p') > c_1 \cdot (p'+k')^{c_2\lfloor\sqrt{2k'}\rfloor}$. Let k' be the smallest integer for which such a pair exists, and let $p' = p'(k')$ be the smallest integer for which $T(k',p') > c_1 \cdot (p' +' k)^{c_2\lfloor\sqrt{2k'}\rfloor}$. Note that $k' \geq 4$ and $T(\ell,q) \leq c_1 \cdot q^3 \cdot (q+\ell)^{c_2\sqrt{\ell}}$ for all pairs (ℓ,q) with $\ell \leq k'$, $q \leq p'$ and $\ell + q < k' + p'$. Further, it holds $c_2 \geq \frac{c_2}{\alpha_k} + 2 + \log(2c_1')$ for all $k \geq 4$. For $k \geq 4$ we also have $\left\lfloor\sqrt{2k}\right\rfloor < k$. This implies that

$$\begin{aligned}
T(k',p') &\leq 2(k'+p')^{\lfloor\sqrt{2k'}\rfloor} \cdot c_1' \cdot k'^{\lfloor\sqrt{2k'}\rfloor} \cdot T(\left\lfloor\sqrt{2k'}\right\rfloor, p')) \\
&\leq 2(k'+p')^{\lfloor\sqrt{2k'}\rfloor} \cdot c_1' \cdot k'^{\lfloor\sqrt{2k'}\rfloor} \cdot c_1 \cdot (p' + \left\lfloor\sqrt{2k'}\right\rfloor)^{c_2\left\lfloor\sqrt{2\lfloor\sqrt{2k'}\rfloor}\right\rfloor} \\
&\leq 2(k'+p')^{\lfloor\sqrt{2k'}\rfloor} \cdot c_1' \cdot k'^{\lfloor\sqrt{2k'}\rfloor} \cdot c_1 \cdot (p' + \left\lfloor\sqrt{2k'}\right\rfloor)^{c_2\left\lfloor\sqrt{2\lfloor\sqrt{2k'}\rfloor}\right\rfloor} \\
&\leq (2c_1'c_1)(k'+p')^{2\lfloor\sqrt{2k'}\rfloor + c_2\left\lfloor\sqrt{2\lfloor\sqrt{2k'}\rfloor}\right\rfloor} \\
&\leq c_1(k'+p')^{2\lfloor\sqrt{2k'}\rfloor + \frac{c_2}{\alpha_{k'}}\lfloor\sqrt{2k'}\rfloor + \log(2c_1')} \\
&\leq c_1(k'+p')^{(2+\frac{c_2}{\alpha_{k'}} + \log(2c_1'))\lfloor\sqrt{2k'}\rfloor} \\
&\leq c_1(k'+p')^{c_2\lfloor\sqrt{2k'}\rfloor}
\end{aligned}$$

This contradicts the existence of k', and therefore concludes the proof. ∎

Lemma 21. *Suppose that*

$$T(k) \leq 2(2k)^{\lfloor\sqrt{2k}\rfloor} \cdot O\left(k^{\lfloor\sqrt{2k}\rfloor} T(\left\lfloor\sqrt{2k}\right\rfloor)\right)$$

and that $T(\ell) \leq c'\ell^{\lfloor\sqrt{2\ell}\rfloor}$ *for some constant* $c' \in \mathbb{R}$ *and for all* $\ell \leq 4$. *Then there exist constants* $c_1 \geq c', c_2 \geq 1$ *such that for all* $k \in \mathbb{N}$,

$$T(k) \leq c_1 k^{c_2\lfloor\sqrt{2k}\rfloor} \ .$$

Proof. Since $T(k) \leq 2(2k)^{\lfloor\sqrt{2k}\rfloor} \cdot O\left(k^{\lfloor\sqrt{2k}\rfloor} T(\left\lfloor\sqrt{2k}\right\rfloor)\right)$, there exists a constant $c_1' > 0$ such that $T(k) \leq 2(2k)^{\lfloor\sqrt{2k}\rfloor} \cdot c_1'\left(k^{\lfloor\sqrt{2k}\rfloor} T(\left\lfloor\sqrt{2k}\right\rfloor)\right)$. Let $\alpha_k = \frac{\lfloor\sqrt{2k}\rfloor}{\left\lfloor\sqrt{2\lfloor\sqrt{2k}\rfloor}\right\rfloor}$. Then for $k \geq 5$ it holds that $\alpha_k \geq 1.5 > 1$. Let $c_1 = \max\{c_1', c'\}$ and let $c_2 = 12 + 3\log c_1$. Suppose, for sake of contradiction, that the statement of the lemma does not holds for this choice of (c_1, c_2). Then exists a $k' \in \mathbb{N}$ such that $T(k') > c_1 k'^{c_2\lfloor\sqrt{2k'}\rfloor}$. Let k' be the smallest such integer. Note that we must have $k' \geq 4$ and $T(\ell) \leq c_1 \ell^{c_2\lfloor\sqrt{2\ell}\rfloor}$ for all $\ell < k'$. Further,

it holds that $c_2 \geq \frac{c_2}{\alpha'_k} + 4 + \log c_1$ for all $k \geq 5$. For $k \geq 4$ we also have $\left\lfloor \sqrt{2k} \right\rfloor < k$. This implies that

$$\begin{aligned}
T(k') &\leq 2(2k')^{\lfloor\sqrt{2k'}\rfloor} \cdot c'_1 \left(k'^{\lfloor\sqrt{2k'}\rfloor} T(\left\lfloor \sqrt{2k'} \right\rfloor) \right) \\
&\leq 2(2k')^{\lfloor\sqrt{2k'}\rfloor} \cdot c_1 \left(2c_1 k'^{\lfloor\sqrt{2k'}\rfloor} k'^{c_2 \left\lfloor \sqrt{2\lfloor\sqrt{2k'}\rfloor} \right\rfloor} \right) \\
&\leq 2c_1(2k')^{\lfloor\sqrt{2k'}\rfloor} k'^{(\frac{c_2}{\alpha'_k}+1)\lfloor\sqrt{2k'}\rfloor + \log c_1} \\
&\leq 2c_1 k'^{2\lfloor\sqrt{2k'}\rfloor} k'^{(\frac{c_2}{\alpha'_k}+1)\lfloor\sqrt{2k'}\rfloor + \log c_1} \\
&\leq c_1 k'^{\lfloor\sqrt{2k'}\rfloor(\frac{c_2}{\alpha'_k}+3)+\log c_1+\log 2} \\
&\leq c_1 k'^{\lfloor\sqrt{2k'}\rfloor(\frac{c_2}{\alpha'_k}+4+\log c_1)} \\
&\leq c_1 k'^{c_2\lfloor\sqrt{2k'}\rfloor} \ .
\end{aligned}$$

This contradicts the existence of k', and therefore concludes the proof. □

Since $T(k,p) \in O(p^{(k^2)})$, it holds that $O(p^{8\lfloor\sqrt{2k}\rfloor})$. Moreover, as $T(k) = O(1)$ for $k \leq 4$, we conclude that $T(k,p) = (p+k)^{O(\sqrt{k})}$ and $T(k) = k^{O(\sqrt{k})}$. This completes the proof of Theorem 5. □

Next, we will prove the following.

Lemma 22. *For $s_j, j, n \in \mathbb{N}$, $s = \left\lceil \sqrt{\frac{s_j}{\log_j s_j}} \right\rceil$ and $\ell = \left\lceil \sqrt{2(n-s_j)} \right\rceil$ we obtain*

$$T(s_j, n-s_j) = n^{O\left(\sqrt{n-s_j}+\sqrt{\frac{s_j}{\log_j s_j}}\right)} \ .$$

To prove Lemma 22, we first establish another lemma.

Lemma 23. *For $s_j, j, n \in \mathbb{N}$, $s = \left\lceil \sqrt{\frac{s_j}{\log_j s_j}} \right\rceil$ and $\ell = \left\lceil \sqrt{2(n-s_j)} \right\rceil$, it holds*

$$T(s_j, n-s_j) \leq n^{O\left(\sqrt{n-s_j}+\sqrt{\frac{s_j}{\log_j s_j}}\right)} \cdot \left(\mathsf{dom}_{n,n-s_j}(\ell, s) + O(1)\right) \ .$$

Proof. For each parity game G on n nodes and $s_j = s_j(G)$ we construct a binary tree T_G in the following way. The root of T_G is labeled by $(s_j, n-s_j)$. A node labeled by (a,b) with $a > 3$ and $b > 3$ has up to two children: a left child labeled by $(a-1,b)$ or $(a,b-1)$, and possibly a right child labeled by $(a - \sqrt{a \cdot \log_j a}, b)$ or $(a, b-\sqrt{b})$. Each child of a node corresponds to one of the recursive calls. The choice on how we label the children depends on the behavior of the algorithm. We label the children of a node by (a',b') and (a'',b'') such that the recursive calls of the algorithm are to games containing at most a', respectively a'' nodes of out-degree at most j and at most b', respectively b'', nodes of out-degree greater than j. Nodes labeled by (a,b) with $a,b \in \{0,1,2,3\}$ are leaves. A node labeled by (a,b) has a cost of $\left(\mathsf{dom}_{a+b,b}(\sqrt{b}, \sqrt{a \cdot \log_j a}) + O(1)\right)$ associated with it.

It follows from Lemma 18 that the sum of the costs of the nodes of T_G is an upper bound on the run time of **new-win**(G,j). The worst possible sum of the costs of the nodes of T_G we can obtain for some instance G with $s_j = s_j(G)$ and $n = |V|$ therefore is an upper bound of $T(s_j, n-s_j)$. Clearly, the length of every path in T_G from the root to a leaf is at most n. We say that such a path makes a *right turn* when it descends from a node to its right child. We next claim that each such path makes at most $O(\sqrt{n-s_j} + \sqrt{\frac{s_j}{\log_j s_j}})$ right turns. This follows immediately from the observation that the function $f(x) = x - \lceil\sqrt{2x}\rceil$ can be iterated on $n-s_j$ at most $O(\sqrt{n-s_j})$ times before reaching the value of 3 or less and the function $g(x) = x - \lceil\sqrt{x \cdot \log_j x}\rceil$

can be iterated on s_j at most $O(\sqrt{\frac{s_j}{\log_j s_j}})$ times before reaching the value of 3 or less. As each leaf of T_G is determined by the positions of the right turns on the path leading to it from the root, we get that the number of leaves in T_G is at most $n^{O\left(\sqrt{n-s_j}+\sqrt{\frac{s_j}{\log_j s_j}}\right)}$. The total number of nodes in T_G is therefore at most $n^{O\left(\sqrt{n-s_j}+\sqrt{\frac{s_j}{\log_j s_j}}\right)}$. As the cost of each node is at most $\left(\mathsf{dom}_{n,n-s_j}(\ell,s)\right)+O(1)$, we immediately have that

$$T(s_j, n-s_j) \leq n^{O\left(\sqrt{n-s_j}+\sqrt{\frac{s_j}{\log_j s_j}}\right)} \cdot \left(\mathsf{dom}_{n,n-s_j}(\ell,s)+O(1)\right). \qquad \square$$

Together with Lemma 15, we obtain

$$T(s_j, n-s_j) = n^{O\left(\sqrt{n-s_j}+\sqrt{\frac{s_j}{\log_j s_j}}\right)}.$$

This completes the proof of Lemma 22.

Acknowledgements. M. M. thanks László Végh for introducing him to parity games, and the authors of [8] for sending us a preprint.

References

[1] D. Berwanger, A. Dawar, P. Hunter, and S. Kreutzer. DAG-width and parity games. In *Proc. STACS 2006*, volume 3884 of *Lecture Notes Comput. Sci.*, pages 524–536. 2006.

[2] D. Berwanger and E. Grädel. Fixed-point logics and solitaire games. *Theory Comput. Syst.*, 37(6):675–694, 2004.

[3] D. Berwanger, E. Grädel, Ł. Kaiser, and R. Rabinovich. Entanglement and the complexity of directed graphs. *Theoret. Comput. Sci.*, 463:2–25, 2012.

[4] H. Björklund, S. Sandberg, and S. Vorobyov. A discrete subexponential algorithm for parity games. In *Proc. STACS 2003*, volume 2607 of *Lecture Notes Comput. Sci.*, pages 663–674. 2003.

[5] E. A. Emerson and C. S. Jutla. Tree automata, mu-calculus and determinacy. In *Proc. FOCS 1991*, pages 368–377, 1991.

[6] J. Fearnley and O. Lachish. Parity games on graphs with medium tree-width. In *Proc. MFCS 2011*, volume 6907 of *Lecture Notes Comput. Sci.*, pages 303–314. 2011.

[7] J. Fearnley and S. Schewe. Time and parallelizability results for parity games with bounded treewidth. In *Proc. ICALP 2012*, volume 7392 of *Lecture Notes Comput. Sci.*, pages 189–200. 2012.

[8] J. Gajarský, M. Lampis, K. Makino, V. Mitsou, and S. Ordyniak. Parameterized algorithms for parity games. In *Proc. MFCS 2015*, Lecture Notes Comput. Sci., pages 336–347, 2015.

[9] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes Comput. Sci.* Springer, 2002.

[10] M. Jurdziński. Deciding the winner in parity games is in UP ∩ co-UP. *Inform. Process. Lett.*, 68(3):119–124, 1998.

[11] M. Jurdziński. Small progress measures for solving parity games. In *Proc. STACS 2000*, volume 1770 of *Lecture Notes Comput. Sci.*, pages 290–301. 2000.

[12] M. Jurdziński, M. Paterson, and U. Zwick. A deterministic subexponential algorithm for solving parity games. *SIAM J. Comput.*, 38(4):1519–1532, 2008.

[13] V. King, O. Kupferman, and M. Y. Vardi. On the complexity of parity word automata. In *Proc. FOSSACS 2001*, volume 2030 of *Lecture Notes Comput. Sci.*, pages 276–286. 2001.

[14] T. Kloks and H. L. Bodlaender. On the treewidth and pathwidth of permutation graphs, 1992. `http://www.cs.uu.nl/research/techreps/repo/CS-1992/1992-13.pdf`.

[15] R. McNaughton. Infinite games played on finite graphs. *Ann. Pure Appl. Logic*, 65(2):149–184, 1993.

[16] J. Obdržálek. Fast μ-calculus model checking when tree-width is bounded. In *Proc. CAV 2003*, volume 2725 of *Lecture Notes Comput. Sci.*, pages 80–92. 2003.

[17] J. Obdržálek. Clique-width and parity games. In *Proc. CSL 2007*, volume 4646 of *Lecture Notes Comput. Sci.*, pages 54–68. 2007.

[18] S. Schewe. Solving parity games in big steps. In *Proc. FSTTCS 2007*, volume 4855 of *Lecture Notes Comput. Sci.*, pages 449–460. 2007.

[19] C. Stirling. Local model checking games. In *Proc. CONCUR 1995*, volume 962 of *Lecture Notes Comput. Sci.*, pages 1–11. 1995.

[20] J. Vöge and M. Jurdziński. A discrete strategy improvement algorithm for solving parity games. In *Proc. CAV 2000*, volume 1855 of *Lecture Notes Comput. Sci.*, pages 202–215. 2000.

[21] W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoret. Comput. Sci.*, 200(1-2):135–183, 1998.